

REMARKS

Claims 1-23 are pending. Claims 1-22 were rejected. Claims 1, 5, 11, 12, and 14-22 have been amended. New claim 23 has been added.

The Office Action indicated that claims 1-12 and 22 are rejected under 35 USC 103(a) as being allegedly unpatentable over Yang, US Patent Application No. 2002/0114334, in view of Haddock, US Patent No. 6,970,426.

The Office Action indicated that claims 14-21 are rejected under 35 USC 101 because the claimed invention allegedly is directed to non-statutory subject matter.

Applicants believe that all pending claims should be allowed, for the reasons cited below.

Applicants' attorneys thank the Examiner for taking the time to share his thoughts regarding the application and the art relied upon during the interview on June 26, 2007.

Replacement Drawings

The Office Action objected to the drawings as being informal. Replacement Figures 1-7 are submitted herewith.

Claim 14: 35 USC 101

The Office Action indicated that claims 14-21 were rejected under 35 USC 101 because the claimed invention was allegedly directed to non-statutory subject matter. Applicants respectfully traverse this rejection. However, in order to expedite the prosecution of this application, Applicants have amended claim 14 as follows:

14. (currently amended): ~~A computer program embodied in a machine-readable medium~~A computer readable medium encoded with a computer program, the computer program comprising instructions for controlling a network device to perform the following steps:

determining a first time at which a first token bucket controlling a first bandwidth allocation is approximately full of first tokens; and

allocating, after the first time and based on the determining step, excess first tokens to a second token bucket controlling a second bandwidth allocation,

wherein the first token bucket corresponds to a first traffic flow and the second token bucket corresponds to a second traffic flow.

The amendments are based on suggestions provided by the Examiner in the Office Action and during the interview of June 26, 2007. Applicants' attorneys thank the Examiner for his suggestions.

Claims 1-10, 11, 12 and 22: 35 USC 103(a)

The Office Action indicated that claims 1-12 and 22 were rejected under 35 USC 103(a) as being allegedly unpatentable over Yang in view of Haddock. Original claim 1 was as follows:

1. (currently amended): A method of allocating bandwidth in a computer network, the method comprising:

determining a first time at which a first token bucket controlling a first bandwidth allocation is approximately full of first tokens; and

allocating, after the first time, excess first tokens to a second token bucket controlling a second bandwidth allocation.

The Office Action admits that the allocating step is not disclosed in Yang. Page 4 of the Office Action states, "Yang, however, does not specifically disclose allocating excess first tokens to a second token bucket..." To support its rejection, the Office Action relies upon Haddock, asserting that Haddock "discloses allocating excess first tokens to a second token bucket E."

Applicants respectfully traverse this rejection. However, in order to expedite the prosecution of this application, Applicants have amended claim 1 to highlight features not believed to be found in either of the prior art references. Amended claim 1 now recites the claim element, "... wherein the first token bucket [controlling a first bandwidth allocation] corresponds to a first traffic flow and the second token bucket [controlling a second bandwidth allocation] corresponds to a second traffic flow."

The amendment is supported by several drawings and passages in the specification, including Figures 3 and 6 and the corresponding descriptions. (See page 13, line 25 through page 15, line 15 and page 18, line 12 through page 20, line 9.) As discussed during the interview of June 26, 2007, both Figure 3 and Figure 6 depict the allocation of excess tokens from a first token bucket (controlling a first bandwidth allocation for a first traffic flow) to a second token bucket (controlling a second bandwidth allocation for a second traffic flow). Similar amendments were made to independent claims 11, 12, 14 and 22.

New claim 23 recites the specific situation “wherein the first traffic flow corresponds with a first service and the second traffic flow corresponds with a second service.” Claim 23 is supported, e.g., by the examples of Figures 3 and 6, wherein each traffic flow corresponds to a separate service.

During the interview of June 26, 2007, Applicants’ attorneys described one such example with reference to Figure 6. There, token bucket 620 controls the “peak information rate” (“PIR”) 625 for the traffic flow of service type 605, whereas token bucket 645 controls the peak information rate 650 for the traffic flow of service type 610. The specification notes that, in one example, the traffic flow of service 605 may be for mission-critical data, the traffic flow of service 610 may be for video data and the traffic flow of service 615 may be for “best-efforts” data. (Specification at page 18, lines 15-19.)

Applicants’ attorneys also noted that in some circumstances, excess tokens may be allocated from a first token bucket that is controlling a first bandwidth allocation for a first traffic flow (e.g., token bucket 620) to a second token bucket controlling a second bandwidth allocation for a second traffic flow (e.g., token bucket 645). “Excess tokens 640 from PIR token bucket 620 may be donated to other PIR token buckets, preferably when PIR token bucket 620 is full or approximately full.” (Specification at page 18, lines 21-24.) Applicants’ attorneys noted that this is symbolized in Figure 6 by the “pipes” that deliver drips of excess tokens from one token bucket to another, e.g., the pipe that delivers excess tokens 640 from token bucket 620 to token bucket 645.

In the foregoing example, a token bucket that is controlling a peak information rate for mission-critical data (token bucket 620) may allocate excess tokens to another token bucket (token bucket 645) that is controlling a peak information rate for video data.

As understood, Haddock does not teach any such feature. Haddock appears to be one example of prior art that is briefly referenced in the Background section of the present application, wherein one or more token buckets are used to enforce a bandwidth profile for a traffic flow. For example, page 4, lines 1-5 state, “[t]ypically, one or more token buckets are used to enforce each bandwidth profile. A particular traffic flow may then be subjected to a particular bandwidth profile. Thus, ***one or more token buckets are used to enforce a bandwidth profile for each traffic flow.***” (Emphasis added.)

During the interview, there was some discussion regarding whether or not Haddock addresses the allocation of excess tokens between token buckets ***corresponding to different traffic flows***. As discussed below, even if Haddock were understood as teaching the allocation of excess tokens from one bucket to another (and it does not), applicants respectfully submit that Haddock still would teach against claim 1, at least because the two token buckets disclosed in Haddock ***correspond to the same traffic flow***.

By way of background, we note that token buckets are described as follows in Wikipedia (see http://en.wikipedia.org/wiki/Token_bucket):

The token bucket is a control mechanism that dictates when traffic can be transmitted, based on the presence of tokens in the bucket. The token bucket contains tokens, each of which can represent a unit of bytes or a single packet. ***Tokens in the bucket are effectively "cached in" (removed) for the ability to send a packet.*** The network administrator specifies how many tokens are needed to transmit how many bytes; ***when tokens are present, a flow is allowed to transmit traffic. If there are no tokens in the bucket, a flow cannot transmit its packets.***

(Emphasis added.)

According to the above citation, a token bucket dictates when its corresponding flow can transmit packets. If a flow corresponds to a token bucket, then a packet transmitted by the flow may deplete the corresponding bucket. In the same vein, if a flow corresponds to two buckets, then a packet transmitted by the flow may deplete either or both buckets.

Figure 1 of Haddock shows that Haddock discloses a correspondence between a single flow and two token buckets. Incoming packets are introduced in block 120. The flow chart shows that any incoming packet may deplete *either* bucket C (block 130) *or* bucket E (block 140). Column 6, lines 15 and 28 state the following:

...if the packet is green and the number of bytes, B , in the packet is less than the number of tokens T_c in token bucket C at time t , at 130 the packet is marked green and *token bucket C is emptied of B tokens*... If the packet is green or yellow and the number of bytes, B , in the packet is less than the number of tokens T_e in token bucket E at time t , at 140 the packet is declared yellow and *token bucket E is decremented by B tokens*.

(Emphasis added.)

Since Haddock identifies only one source of packets (block 120) and since any such packet may deplete either bucket C or bucket E, buckets C and E correspond to the same source of packets, i.e., the same flow. Haddock does not refer to flows or sources of packets that always deplete only one bucket but not the other. Moreover, column 1, line 6 of Haddock describes its invention as a “device metering a received data stream.” Applicants’ attorneys have found no references to multiple data streams or flows in Haddock.

Embodiments of the claimed invention, by contrast, enable the allocation of excess tokens between token buckets corresponding to *different* traffic flows. Figure 3 of the specification displays an embodiment with token buckets 305, 325 and 355. Each token bucket in Figure 3 corresponds to a traffic flow of a corresponding service, denoted by arrows T1, T2 and T3. The specification of the present application notes:

Tokens 301 flow into token bucket 305 at a rate R_1 are, at times, consumed by traffic T1 of the first type of service. However, there will be times during which there is little or no traffic T1. At such times, excess tokens 310 will be diverted or “donated” to another token bucket when token bucket 305 is full at an approximate average rate of $(R_1 - T_1)$... Likewise, tokens 321 flow into token bucket 325 at a rate R_2 and are, at times, consumed by traffic T2 of the second type of service. However, at times when there is little or no traffic T2, excess

tokens 340 will be donated to another token bucket. In this example, excess tokens 340 are donated to token bucket 355 and made available for use by a third type of service.

(Page 13, lines 3-9 and 16-22.)

The discussed distinction between the claimed invention and Haddock is further supported by Haddock's use of the terms CIR and PIR. Page 1, line 30 of the specification describes CIR and PIR as follows:

Often the usage of network resources by a customer is metered or limited by measuring the amount of bandwidth consumed by that customer's traffic as it is submitted to the service provider's network. The term "committed information rate" ("CIR") is often used to denote the maximum amount of bandwidth able to be consumed by a particular customer as measured over some period of time. To allow for fluctuations in the flow of traffic submitted to the network, it is also possible to define an allowable "burstiness." This is often called the "peak information rate" ("PIR").

According to the citation, CIR and PIR are recognized terms that are often used together in connection with a single customer's traffic flow.

Column 4, line 13 of Haddock makes use of the terms CIR and PIR as follows:

The rate color marker comprises a meter that... deposits tokens into token bucket C at rate CIR until token bucket C is full, and deposits tokens into token bucket E at rate CIR if token bucket C is full. Additionally, a separate source of tokens is provided for token bucket E... The combination of the two sources of tokens for deposit in token bucket E, one source deposited at rate CIR if token bucket C is full and the other source deposited at EIR, provides an effective token deposit rate of PIR.

(Emphasis added.)

As noted earlier, CIR and PIR often are used in connection with a single flow. The algorithm disclosed in Haddock does not mention any buckets other than C and E. Haddock, by

associating the term CIR with bucket C and PIR with bucket E, indicates that buckets C and E also correspond to a single flow.

Because Haddock only describes a single traffic flow and two buckets, it cannot reasonably be understood to teach an allocation of excess tokens between different buckets corresponding to different traffic flows. Unlike the present application, Haddock does not suggest such an allocation, much less provide any techniques for performing such an allocation.

CONCLUSION

Haddock does not teach or suggest the allocation of tokens from a first token bucket corresponding to a first traffic flow to a second token bucket corresponding to a second traffic flow. Therefore, Applicants request that the rejection of claim 1 be withdrawn. Since amended independent claims 11, 12, 14 and 22 have similar claim elements, applicants respectfully request that these claims be allowed as well.

Accordingly, Applicants believe that all pending claims are allowable and respectfully request a Notice of Allowance for this application from the Examiner. Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

If any fees are due in connection with the filing of this Amendment, the Commissioner is authorized to deduct such fees from the undersigned's Deposit Account No. 50-0388 (Order No. **CISCP344**).

Respectfully submitted,
BEYER WEAVER LLP
/Roger S. Sampson/

Roger S. Sampson
Reg. No. 44314

P.O. Box 70250
Oakland, CA 94612-0250
(510) 663-1100